Doppler Spectrum Extraction of Planetary Radar Data Using Computer FFT and Integration

D. E. Wallis
Communications Systems Research Section

A computer program to extract integrated doppler spectra of planetary radar returns is described. The program input data are discrete, complex-valued time series obtained for each of 64 range gates. The data record is subdivided into successive subseries of 256 time-points, and each 256-point time series is discrete-Fourier-transformed using the fast Fourier transform (FFT). Power spectra are obtained by magnitude-squaring. The power spectra, containing positive- and negative-frequency doppler components, are assumed statistically stationary during receive-runs of 10 to 15 minutes duration, and so may be (SNR)-enhanced by integration (summation) of successive 256-point power spectra. Program output is a highly compressed power spectrum tape containing 64 integrated spectra, together with identification and parameter values, for all receive runs made during each day of a total of six days of radar observations of the planet Mercury in calendar year 1974. Brief discussions of the on-site data-taking process and the doppler interpretation of the discrete Fourier transform are included.

I. Introduction

This report describes a computer program to extract integrated doppler power spectra of planetary radar returns. The description of the on-site data-taking process and data-reduction requirements in Sect. II is used to introduce the detailed descriptions of data-tape formats in Sect. III and the computational flowchart in Sect. IV. To anticipate and answer some questions of users of the program, a brief discussion of the doppler interpretation of the Discrete Fourier Transform (DFT) is included in Sect. V.

II. Background: The On-Site Data-Taking Process

In the radar observations of the planet Mercury in summer 1974, digital complex-valued (quadrature-pair) amplitude samples of approximately 15-bit maximum resolution are obtained from the radar receiver-detector system. The detector system has 64 range gates, 32 of which are assigned to Coder 1, with the remainder assigned to Coder 2. The operations of Coder 1 and Coder 2 are interleaved, so that Coder 1 produces data for even ranges $0,2,4,\cdots,62$ and Coder 2 produces data for odd ranges $1,3,5,\cdots,63$, where $0,1,2,3,\cdots,62,63$ is the ordering of range itself. At each range assignment, the coders output their complex-valued amplitude samples to the on-site XDS 930 computer, which merges the 128 numbers from the coders with serial-number, time-of-day, and various experimental parameters to form a 192-word record to be written on standard 7-track tape at 800 bits/inch. Each 192-word record, then, contains 64 complex-valued amplitude samples, the set of which may be regarded, for purposes of power-spectrum computation, as representing a vector-valued sample at a single point in time. Two tape drives are operational for the on-site data recording process, so that one drive unit is available on line to the computer while the other unit is being rewound and reloaded. In a given day's observations, 18 to 19 reels of tape are written, and the succession of 192-word records is continuous from each tape reel to the next, except that a 4K-word core dump is always the first record written at the beginning of each reel used. The condition for continuing the data recording onto the available alternate tape drive is, effectively: If not end-of-tape, write to present unit; if end-of-tape, switch units, write a core-dump record, and resume data recording without loss of continuity. Because of variable distances to the planet under observation, the number of records written during each receive-run (duration of approximately one round-trip-light-time) varies. To maximize utilization of recording tape, the recording process is suspended in mid-tape at the end of a receive run, so that recording may resume without loss of footage as data from the next receive run become ready to record. Each 192-word, time-point record contains a serial number which serves, along with a notation of time-of-day for each record, to ensure that data for truly successive points in time can be identified in the subsequent data reduction. The boxes of tapes for a given day's observations are accompanied by a computer typeout of the serial numbers marking the ends of the successive receive-runs, together with other station log information not used in the present program.

A. Requirements for the Data Reduction Program

The computer program to be described in the following must, of course, contain program logic to read data tapes organized as described above. Additionally, the program must anticipate and guard against possible errors in the on-site data-taking process which, due to heavy real-time computational demands, cannot be corrected on-line without incurring other unacceptable losses. A further requirement on the present program is to produce intercomputer compatible output, so that the power spectrum tapes to be produced are suitable for dissemination to users demanding high confidence in data validity and processing integrity.

Embedded within the above requirements, are the more purely mathematical operations of the discrete Fourier transform, spectrum computation, and spectrum integration. The Fourier transform is accomplished by the now traditional FFT technique, and program requirements reduce to ensuring that computations are correctly scaled, free from overflows, that spectral extraction is correct, and that output formats are as specified.

III. Input and Output Tape and Record Formats

Input to the computer program comprises time tapes generated at the station, plus certain entries obtained from the radar log, which are typed into the computer at processing time by the computer operator. Output from the computer program is an intercomputer compatible spectrum tape, incidentally accompanied by a computer typeout of important processing milestones, plus a record of occurrences of computational overflow. The input and output tapes are 7-track, 800 bits/in. and their formats are given in Tables 1–4. Explanatory remarks are included with the figures. Section II (q.v.) may be consulted in connection with the time tape formats. The content of the spectrum tape header records is for the most part copied from certain portions of the time tape data records, but the content of the spectrum tape power spectrum records is derived by Discrete Fourier Transform (DFT) and magnitude-squaring from the time-tape time series.

IV. Description of Program Flowchart

At the start of the program (Fig. 1), the table of values of the complex exponentials required by the DFT/FFT is computed. (See Sect. V for a discussion of the exponential table). Also computed at the start of program is a table of indirect address keys designed to manipulate any 256-point time series of normal sequential ordering into the so-called "shuffled" or "sorted" order required by the particular canonical form of FFT used.

The program asks for input of date parameters required to be included in the spectrum tape output from the program. Also, in advance of each receive run to be processed, the program asks for the record serial number delimiting the last record of the run. The program tests for end-of-tape prior to reading each data record and calls for the next tape in sequence to be mounted. The core dump record, which is the first record written on each new reel of tape, is always skipped. Since the integration of power spectra is always over all spectra computed in the course of a given receive run, the integration accumulator is cleared to zero prior to processing data for the run.

The program's logic effectively makes the reading of successive records continuous in spite of the reel-to-reel transition, and does not "lose" any records either reel-to-reel or run-to-run. Processing of a given receive run terminates when a record of greater serial number than that for end-of-run is encountered. This record is retained and used as the first record of the next receive run. The requirement that only continuous, successive time-records be FFT'ed is enforced by checking the incrementation of record serial numbers and also checking the incrementation of the recorded time-of-day words on the successive records.

Data from the time tapes are read in 256-record blocks in order to gather the time series to be transformed. Somewhat unfortunately, however, 256-point complex-valued time series are obtained for all 64 range gates at once. Since a total of 32,768 data words are obtained, it was found necessary to use drum storage (RAD) to hold the full data set. The program, therefore, organizes the reading of a 256-record block as eight (8) 32-record blocks, and exploits the fact that each 32-record block (4096 words) conveniently fills one band of the drum, thereby tending to minimize drum loading time. The 4096-word record written to the drum band is derived by reading 32 records from tape to core memory in such a way that the 32 successive complex numbers for each range gate versus time are grouped contiguously.

The data for each range gate are made available to the FFT computations by reading selectively from the 8 drum bands containing the total 256-record block and loading these data to a scratch pad area of core memory. The time-sequential order of the data is altered to the "shuffled" order used by the transform. A conventional complex FFT is done for each range-gate's data, with monitoring of any computational overflows which may occur anywhere in the processing of the 32,768-word block. Since the first receive run of a given day's records is actually a delay calibration of the transmitter-receiver system, rather than a planetary return, the somewhat larger measurement values of the delay calibration are handled by right shifting all data 2 bits, with subsequent left shift correction in the computed results. The 256-point (512-word) complex-valued transform computed by the FFT for each range gate is magnitude-squared to produce a 256-word single-precision power spectrum for the gate. The power spectrum is accumulated in a double-precision accumulator assigned to each gate number.

When all full 256-record blocks of the receive run have been processed (with any remaining number of records less than 256 being ignored), the integrated spectrum for each gate is normalized by left-shifting until the largest spectral element is scaled to maximum positive magnitude. The number of shifts required for the maximum element is then applied to all elements of the spectrum, thereby yielding a normalized spectrum for the gate. The number of shifts required to normalize each spectrum is incorporated into a computed scale factor, whereby the normalized spectrum can later be un-normalized, if desired, by the user.

When processing of data for each receive run is finished, the 65 records for the computed run are written on the output spectrum tape. The successive 65-record blocks comprise the output for each day's records. In the Mercury data processed, data for six days records appear on the spectrum tape.

V. Doppler Interpretation of the Discrete Fourier Transform

The following necessarily brief discussion is intended to address various user questions which have arisen in connection with the content and mathematical interpretation of doppler spectra generated by the present computer program.

Suppose that a discrete time series of N complex-valued detector-output amplitude samples $\{Z_i\}$, $0 \leq i \leq N-1$, has been obtained at a uniform sampling rate. Suppose further that the discrete samples are actually samples of a continuous waveform representable as a sum of harmonics in the usual sense of Fourier analysis, in which the frequencies of the assumed harmonics are expressible as integer multiples of a lowest- or fundamental-frequency sinusoid. With N discrete samples available, it is known that N harmonics are sufficient to fit the time series exactly. The N assumed harmonics comprise a dc (constant) waveform, plus $N - 1$ sinusoids whose frequencies are integer multiples of a sinusoid having one period of oscillation in the time required to gather the N samples. In planetary radar, the amplitudes of all the harmonics are of interest, and these are computed by the Discrete Fourier Transform (DFT). The DFT is a computation which may be indicated by (but is not usually computed by) the following symmetric $N \times N$ matrix times N-vector multiplication:

$$\begin{bmatrix} s_0 \\ s_1 \\ \vdots \\ s_{N-1} \end{bmatrix} = \frac{1}{N} \begin{bmatrix} \\ \exp\left(-j\frac{2\pi}{N}nm\right) \\ \\ \end{bmatrix} \begin{bmatrix} Z_0 \\ Z_1 \\ \vdots \\ Z_{N-1} \end{bmatrix}$$

N = dimension of transform

$0 \leq n < N - 1$ column index

$0 \leq m < N - 1$ row index

j = imaginary unit

Recalling that the determination of the harmonic amplitudes in a conventional Fourier analysis involved multiplication of the given time function by one of the assumed harmonic functions and integrating, it will be evident that the DFT involves an analogous procedure. The analogue may be seen by examination of any of the row-times-column products in the multiplication: Each row of the DFT matrix is simply a tabulation of the complex exponential (sinusoid) for integer frequency m (row index) versus time n (column index). Thus, the complex exponential for frequency m is being multiplied by the time series $\{Z_i\}$, with summation (integration) of this product with respect to time to produce the harmonic amplitude s_m for frequency m. This process, for a given row (frequency), will extract the value (Fourier amplitude coefficient) of *one* of the unknown amplitudes. If the time series $\{Z_i\}$ is assumed to be represented as a sum of harmonics

$$Z_m = \sum_{n=0}^{N-1} c_n \exp\left(+j\frac{2\pi}{N}nm\right) \qquad \begin{array}{l} 0 \leq m \leq N-1 \text{ (time)} \\ 0 \leq n \leq N-1 \text{ (frequency)} \end{array}$$

where n may be regarded as the harmonic number and m as the time argument, then the product of the kth row of the DFT matrix times the time-series,

$$s_k = \frac{1}{N}\sum_{m=0}^{N-1} \exp\left(-j\frac{2\pi}{N}km\right)\left\{\sum_{n=0}^{N-1} c_n \exp\left(+j\frac{2\pi}{N}nm\right)\right\}$$

is well known to reduce to the result $s_k = c_k$.

Thus, the kth row of the DFT matrix, which is a negative-exponent, complex exponential of the form

$$\exp\left[-j\frac{2\pi}{N}km\right]$$

extracts from the time series the Fourier amplitude coefficient of any component, or *harmonic content*, having the positive-exponent, complex exponential form

$$\exp\left[+j\frac{2\pi}{N}km\right]$$

The DFT, then, is an operation which computes all of the Fourier amplitude coefficients at once. This set, represented as the column vector $\{s_i\}$, $0 \leq i \leq N - 1$, is called the discrete Fourier transform, or discrete Fourier amplitude spectrum, of the given time series. The set of N rows of the DFT matrix is sufficient, i.e., *complete*, so that not only does every time series have a computable transform, but also that from the transform the time series may be reconstructed exactly. Thus, the transform is invertible.

Before proceeding to a doppler interpretation of the DFT, it is well to mention that the scalar factor $1/N$ which multiplies all elements of the matrix-vector product is usually dropped in discussions of the computational implementation of the DFT, e.g., by the so-called Fast Fourier Transform (FFT). Indeed, the factor $1/N$ is not included in the calculations accomplished by the present computer program. The factor $1/N$ is required, however, for certain mathematical purposes other than relative spectrum computation, such as expressing *Parseval's equation* (conservation of power), and for making the DFT matrix *unitary*, if such matters are of interest.

Returning now to the interpretation of the DFT, examination of the rows of the matrix reveals that the (positive-exponent) 0th,1st,2nd, $\cdots$, $(N-1)$st frequency component amplitudes in the time series will be extracted by the corresponding rows of the transform matrix. This observation is correct, but the following additional remarks bear on the interpretation: Suppose $N = 256$, for example. Then the highest-frequency harmonic in the time-series is understood to be the 255th. The 255th harmonic is visualized as a sinusoidal waveform having 255 periods of oscillation in the period of the fundamental. If this 255th harmonic were tabulated at the Nth time-points (and no others) the tabulation happens not to show the fine structure expected to reflect 255 periods of oscillation but, rather, would seem to show only *one* period of oscillation. Thus, the *apparent* frequency of the 255th harmonic in the tabulation is actually that of the fundamental. More precisely, since complex exponentials are involved, the tabulation for harmonic number 255 coincides with the tabulation for harmonic number -1, if harmonic number 1 is taken as the fundamental. These remarks will introduce the applicability of the identity

$$\exp\left(-j\frac{2\pi}{N}nm\right) \equiv \exp\left[-j\frac{2\pi}{N}(nm \bmod N)\right]$$

and, more specifically,

$$\exp\left(-j\frac{2\pi}{N}nm\right) \equiv \exp\left(-j\frac{2\pi}{N}[-N-n)]m\right)$$

From the latter identity, for $0 \leq n,m < N$, it may be seen that postulating high positive harmonic numbers, such as 255, is mathematically indistinguishable from postulating low *negative* harmonic numbers, such as -1. The *completeness* of the transform with respect to postulated frequency components at the harmonic numbers $0,1,\cdots,(N-1)$ is retained if arbitrary (and possibly noncontiguous) harmonic numbers outside this range are postulated, provided that the set of the chosen harmonic numbers is congruent to $0,1,\cdots,(N-1)$ *modulo* N. Thus instead of postulating harmonic numbers $0,1,2,\cdots,255$, it would be entirely equivalent, for example, to postulate that the frequencies $+128$, -127, $-126,\cdots,-1,0,1,2,\cdots,$ $126,127$ are contained in a 256-point time series. For doppler amplitude samples obtained from planetary radar, for example, such a choice seems natural: on physical grounds, it is known *a priori* that (a) the doppler return does in fact contain frequency components above, below, and at the carrier frequency, and (b) all frequency components lie in the immediate neighborhood of the carrier frequency. Now, heterodyning of the doppler return to dc (zero frequency) would be expected to produce positive- and negative-frequency components with respect to dc. Therefore, if the function of the DFT is to extract a spectrum of this form, the *user* is postulating a set of harmonic numbers such as $+128$, -127, $-126,\cdots,$ $-1,0,1,2,\cdots,126,127$ rather than any of the other choices he could have made congruent modulo 256. The burden of justifying an assignment of the spectrum values computed by the DFT to any postulated set of frequencies, then, necessarily rests with the user.

The present computer program accommodates the $+128$, -127, $-126,\cdots,-1,0,1,2,\cdots,125,126,127$ assignment by computing the DFT *without* regard to this assignment, and then permuting the order of the elements so as to reflect the user's view. The permutation may be indicated schematically as:

$$\text{from DFT}\begin{bmatrix} s_0 \\ s_1 \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ s_{255} \end{bmatrix} \quad \text{to user}\begin{bmatrix} s_{128} \\ s_{129} \\ \vdots \\ s_{254} \\ s_{255} \\ s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{126} \\ s_{127} \end{bmatrix} \left.\begin{array}{l} \leftarrow \text{frequency having no} \pm \text{sense} \\ \leftarrow \text{most negative frequency} \\ \\ \\ \\ \leftarrow \text{dc} \\ \\ \\ \\ \\ \leftarrow \text{most positive frequency} \end{array}\right\} \begin{array}{c}\text{symmetrical} \\ \text{frequency spectrum}\end{array}$$

With an even number of points assumed in the time series (256 in the present discussion) the $+128$th harmonic number gives a waveform identical to that for the -128th ($+128 \equiv -128 \bmod 256$). Since only one of these could be computed (either, with the same computational result), there seems to be no meaningful interpretation of s_{128} as a doppler frequency amplitude component. Interpretation is not, however, a part of the computational requirements. Thus, s_{128} is computed for the sake of completeness of the transform, and is furnished to the user along with the other spectral elements.

Intentionally omitted from the foregoing discussion was elaboration of the fact that a *power* spectrum is the computational result furnished to the user. A power spectrum is obtained from the amplitude spectrum (computed by DFT) by magnitude-squaring the elements, and summing all spectra computed in successive 256-point blocks making up a given receive run. Thus, the user gets:

$$\sum_{\substack{\text{all trans-}\\\text{forms in a}\\\text{receive-run}}} \begin{bmatrix} |s_{128}|^2 \\ |s_{129}|^2 \\ \vdots \\ |s_0|^2 \\ \vdots \\ |s_{126}|^2 \\ |s_{127}|^2 \end{bmatrix}$$

By the statistical stationarity assumed for the spectra of successive 256-point time series, the computed spectra are simply SNR-enhanced by the summation, i.e., an SNR-improved estimate of the doppler spectrum is obtained. The magnitude-squaring operation wipes out the phase information for each frequency component, and preserves only a power measure of each component.

Table 1. Time tape record format and contents

Time tape record format	
BOT marker	
Core dump record: 4096 words, dumped from memory locations 0-4095_{10}.	
Data record: 192 words	all data records formatted identically
Data record: 192 words	
⋮	
Data record: 192 words	
EOT marker	

Time tape 192-word data record contents

Word No. (1-192)				Range order of complex pairs
1	Record serial number (binary integer)			
2	Time of day in seconds (binary integer)			
3	Parameters			
4				
⋮				
63				
64				
65	Real sample	Gate 0	Coder 1	0
66	Imag sample			
67	Real sample	Gate 1		2
68	Imag sample			
⋮		⋮		⋮
127	Real sample	Gate 31		62
128	Imag sample			
129	Real sample	Gate 0	Coder 2	1
130	Imag sample			
131	Real sample	Gate 1		3
132	Imag sample			
⋮				⋮
191	Real sample	Gate 31		63
192	Imag sample			

Table 2. Spectrum tape record format

BOT		
Long erasure		
EOF		
Header record	Receive run 0	1st Days records
Spectrum record, range 0		
Spectrum record, range 1		
⋮		
Spectrum record, range 63		
Header record	Receive run 1	
Spectrum record, range 0		
Spectrum record, range 1		
⋮		
Spectrum record, range 63		
⋮		
Header record	Last receive run of day	
Spectrum		
EOF		
⋮		
EOF		
Header record	Receive run 0	6th Days records
Spectrum record, range 0		
Spectrum record, range 1		
⋮		
Spectrum record, range 63		
⋮	Receive run 1	
⋮	Last receive run of day	
EOF		
EOF (extra EOF)		

Table 3. Spectrum tape 197-word header record contents

Word No. (1-197)	
1	ID initially 00000000, increments 1000_8 after each receive run of day
2	No. of words in record −2, = 195_{10}, binary integer
3	Record type = 5 if Run 0, otherwise = 4, binary integer
4	Date record written (processed), 6-digit BCD packed in single 24-bit word, YYMMDD
5	Date data taken (from radar log), 6-digit BCD, single 24-bit word, YYMMDD
6 7 . . 68 69	64-word record, exact copy of time tape data record header, for the *first* time point of the *first* continuous 256-record time-tape data block successfully transformed in the receive run
70 71 . . 132 133	64-word record, exact copy of time tape data record header, for the *last* time point of the *last* full continuous 256-record time-tape data block successfully transformed in the receive run
134 135 . . 196 197	63-word record, listing the record serial numbers of the first records of all of the full, continuous 256-record blocks whose power spectra contribute to the accumulation of spectra for the receive run.

Table 4. Spectrum tape 262-word power spectrum record contents

Word No. (1-262)		
1	ID 1-word field has the 8-digit octal form rrrrrggg. Field rrrrr is initially zero; and increments by 1 after each receive run of day. Field ggg is the gate number for which the spectrum was computed. Field ggg is of the form abb: a = 1 or 2 (for Coder number 1 or 2, respectively), and bb is the gate number $0\text{-}31_8$.	
2	No. of words in record-2, = 260_{10}, binary integer	
3	Record type = 3, binary integer	
4	Date record written (processed), 6-digit BCD packed in single 24-bit word, YYMMDD	
5	Date data taken (from radar log), 6-digit BCD packed in single 24-bit word, YYMMDD	
6	SF (scale factor). A positive binary integer giving the number of right shifts required, for all words in the power spectrum, to restore the absolute scaling of the spectrum. The absolute scaling is that which obtains when integer time series are Fourier-transformed using the conventional FFT (with factor 1/256 omitted), with squaring and accumulation of spectra for all 256-record blocks in a given receive run.	
7	Power spectrum point, frequency ±128	256-point power spectrum for a given range gate
8	Power spectrum point, frequency −127	
9	Power spectrum point, frequency −126	
.		
.		
135	Power spectrum point, frequency 0	
.		
.		
261	Power spectrum point, frequency +126	
262	Power spectrum point, frequency +127	

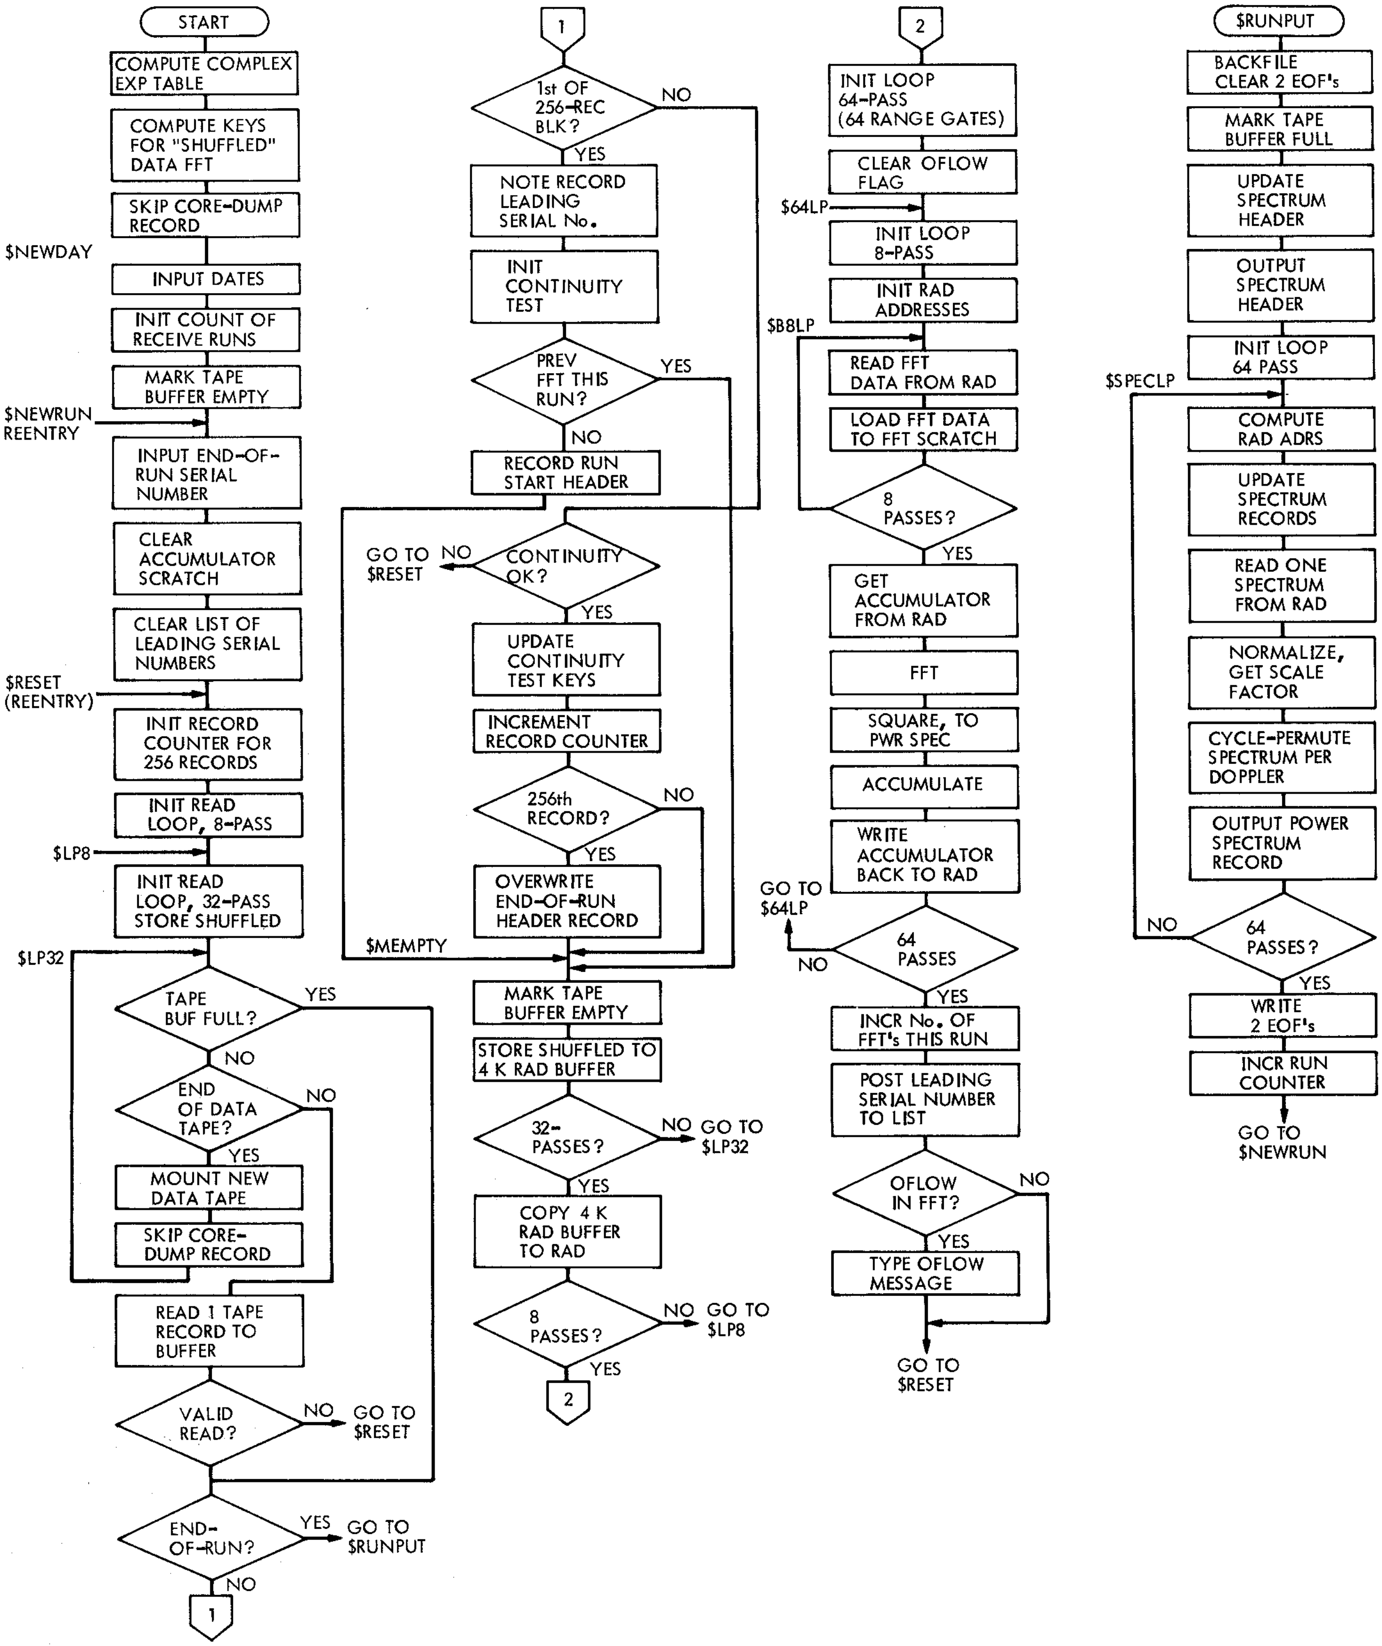

Fig. 1. Computer program flow chart